A Fingerprint and Timing-based Snooping Attack on Residential Sensor Systems

Vijay Srinivasan, John Stankovic and Kamin Whitehouse
Department of Computer Science
University of Virginia
{vs8h,whitehouse,stankovic}@cs.virginia.edu

Abstract

We demonstrate a Fingerprint and Timing-based Snooping (FATS) attack in which private resident information is inferred from the wireless transmissions of sensors in a home, even when all of the transmissions are encrypted. This attack can already be carried out on millions of homes today, and will become more important as ubiquitous computing environments become more prevalent. We address this problem using a suite of privacy solutions and argue that privacy of monitored individuals must be addressed in all design layers in future cyber-physical systems.

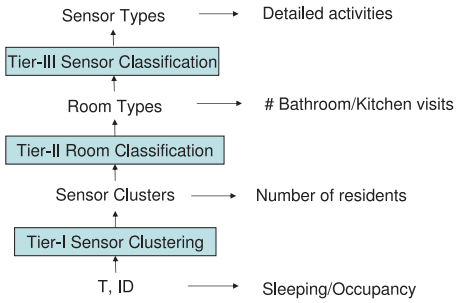

Figure 1. Inference Architecture

It is envisioned that many cyber physical systems of the future will be people-centric, and will consist of sensor-actuator systems deployed in homes and businesses to interact with people in interesting and useful ways. As a result of the *sheer* ubiquitousness of these systems, private information pertaining to individuals can sometimes be inferred by adversaries even if traditional privacy policies and mechanisms are in place. In this work, we propose a new type of privacy attack on *residential wireless cyber physical systems*, typically used in the home automation/security domain and in the medical domain for elderly monitoring. We demonstrate that an adversary can infer detailed resident activity information including sleep duration, bathroom and kitchen visits, and more fine-grained activities *using just the wireless transmission timing and fingerprint information*, where *wireless fingerprinting* techniques use RF features to differentiate radio sources. We do not assume any knowledge about radio message content, type and distribution of devices inside the home. We term this attack, the FATS attack(Fingerprint and Timing Based Snoop Attack).

Our proposed privacy attack comes under the class of side-channel privacy attacks, where the adversary infers private information in spite of encryption by observing how the physical implementation of the system operates. Examples of such attacks include timing and power based attacks on encryption devices [1] and traffic analysis attacks in wired and wireless network systems. In our work, the adversary uses just transmission timing and wireless fingerprint information to infer surprisingly detailed private resident information in spite of encryption. The use of wireless fingerprint information [2] to pose the FATS attack is novel, since fingeprinting has traditionally been used to enhance privacy by providing hardware-based authentication. To demonstrate and evaluate the FATS attack, we deployed wireless X10 sensors for at least seven days in up to four single person homes and four multi-person homes to collect real sensor data relating to resident activity. Around 15 sensors were deployed in each home including motion sensors in each room and contact sensors on objects of everyday use such as the microwave, flush, refrigerator and the front door. Given this real data, we propose a tiered inference algorithm the adversary uses to infer detailed activity information using just transmission timing and the wireless fingerprint. An overview of the inference architecture is shown in Figure 1.

We now provide a brief description of the inference techniques used in each tier. The initial input to the architecture is a set of (TimeStamp,Fingerprint) pairs. *Tier-0* simply infers home occupancy and sleep duration by observing long silence periods during the day and night respectively. In *Tier-I*, we cluster the devices using k-means into spatial groups corresponding to rooms, under the assumption that devices in the same room fire at similar times. In *Tier-II*, we assign a room label to each spatial cluster from a pre-defined set of commonly occuring labels such as bath-

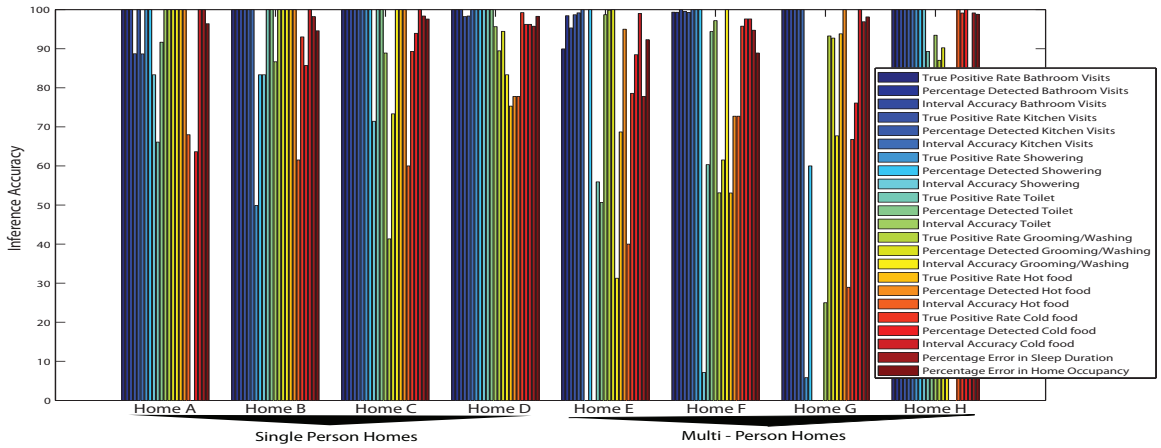

Figure 2. Overall Inference Results in single and multi-person homes

room, bedroom, kitchen etc, using supervised learning on temporal features of room usage; we use a novel weighted matching algorithm on a distance matrix computed using Euclidean distances between the test spatial cluster features and the Gaussian mean features in the training model. In Tier-II, we infer the timing and duration of bathroom and kitchen visits given our knowledge of room labels. In *Tier-III*, we classify radio sources to one of the typical bathroom or kitchen sensors such as refridgerator or shower sensor using the room information in Tier-II. We use an LDA(Linear Discriminant Analysis) classifier trained on temporal features of individual radio sources to infer sensor types. We then use the sensor type information to classify temporal activity clusters in each room, again using LDA classification, as one of showering, toileting, washing or coooking hot/cold food. We show the inference results for all private variables in **Figure 2** in the best case across the four single and multi-person homes. We evaluate adversarial inference accuracy by measuring *true positive rate, percentage detected and interval errors* for bathroom and kitchen visits and the more detailed resident activities and *percentage error* for sleep and home occupancy duration. As seen, the inference accuracy is high, around 90-100% across all homes for bathroom and kitchen visits, while detailed activity inference is much more accurate than possible from random guessing, especially in single-person homes. This graph illustrates the importance of the FATS attack and the need to address it in future cyber physical systems.

We propose and evaluate a suite of privacy solutions with different tradeoffs in our work a) Faraday cages to prevent wireless snooping, b) Periodic transmissions for all devices c) Fingerprint masking to defeat wireless fingerprinting d) Random delays on radio messages and e) Introducing fake data. Random delays increase inference error at low energy cost, but do not *guarantee* complete privacy, while periodic transmissions guarantee privacy, but at a higher energy cost; neither solution can handle real-time requirements. Fingerprint masking and Faraday cages involve possibly expensive changes to existing hardware. Thus, each solution has different tradeoffs in terms of cost, power and privacy guarantees.

In general, cyber physical systems of the future should be designed in such a way that private information cannot be inferred from information that the system makes public. Privacy violations through wireless snooping and fingerprinting are one instance of such an inference, and need to be investigated further before large scale ubiquitous system deployments become a reality. Further research in this area would involve contributions from varied disciplines, from investigating possible fingerprint masking techniques at the hardware layer to devising robust adversarial inference techniques for other interesting attack scenarios. Private variable inference can also occur by mining *aggregate* data pertaining to many individuals made public by large scale cyber physical systems, unless privacy mechanisms are employed to prevent such an inference. Thus, privacy issues must be considered at all design layers in future cyber physical systems, from the hardware layer to the data publishing layer, to ensure privacy of individuals and businesses.

References

[1] H. Bar-El. *Introduction to Side Channel Attacks*. http://www.hbarel.com/publications.htm.

[2] J. Hall, M. Barbeau, and E. Kranakis. Detecting rogue devices in bluetooth networks using radio frequency fingerprinting. In *IASTED International Conference on Communications and Computer Networks*, October 2006.